

Extracting Unsatisfiable Cores for LTL via Temporal Resolution

(full version; r352 December 18, 2012)

Viktor Schuppan (Email: Viktor.Schuppan@gmx.de)

Abstract. Unsatisfiable cores (UCs) are a well established means for debugging in a declarative setting. Still, tools that perform automated extraction of UCs for LTL are scarce. Using resolution graphs to extract UCs is common in many domains. In this paper we show how to construct resolution graphs for temporal resolution as implemented in the temporal resolution-based solver TRP++ and how to use them to extract UCs for propositional LTL. We implement our method in TRP++, and we experimentally evaluate it. Source code of our tool is available.

1 Introduction

Motivation Unsatisfiable cores (UCs) are well established in formal verification, with important uses being debugging in a declarative setting (e.g., [BDTW93]) and avoiding the exploration of parts of a search space that can be known not to contain a solution for reasons “equivalent” to the reasons for previous failures (e.g., [CTVW03]).

LTL (e.g., [Pnu77, Eme90]) and its relatives are important specification languages for reactive systems (e.g., [EF06]). Experience in verification (e.g., [BBDER01, Kup06]) as well as in synthesis (e.g., [BGJ⁺07]) has lead to specifications themselves becoming objects of analysis. Consider two ways to examine a specification ϕ in LTL [PSC⁺06]. First, one can ask whether a certain scenario ϕ' , also given as an LTL formula, is permitted by ϕ . That is the case iff $\phi \wedge \phi'$ is satisfiable. Second, one can check whether ϕ ensures a certain property ϕ'' given in LTL. ϕ'' holds in ϕ iff $\phi \wedge \neg\phi''$ is unsatisfiable. In the first case, if the scenario turns out not to be permitted by the specification, a UC can help to understand which parts of the specification and the scenario are responsible for that. In the second case a UC can show which parts of the specification imply the property. Moreover, if there are parts of the property that are not part of the UC, then those parts of the property could be strengthened without invalidating the property in the specification; i.e., the property is vacuously satisfied (e.g., [BBDER01, KV03, AFF⁺03, FKSfV08, Kup06]). Despite this relevance interest in UCs for LTL has been somewhat limited (e.g., [CRST07, Sch12b, HH11]). In particular, publicly available tools that automatically extract fine-grained UCs for propositional LTL are scarce.

Extracting UCs is often possible using any solver for the logic under consideration by weakening subformulas one by one and using the solver to test

whether the weakened formula is still unsatisfiable (e.g., [Sil10]). While that is simple to implement, repeated testing for preservation of unsatisfiability may impose a significant run time burden. Hence, it is interesting to investigate methods to extract UCs from a single run of a solver. Extracting UCs from resolution graphs is common in propositional SAT (e.g., [GN03,ZM03a]). A resolution method (e.g., [BG01,Rob65]) for LTL, temporal resolution (TR), was suggested by Fisher [Fis91,FDP01] and implemented in TRP++ [HK04,HK03,trp].

Contributions In this paper we make the following contributions. 1. We construct resolution graphs for TR for propositional LTL as implemented in TRP++ [HK04,HK03,trp], and we show how to use them to extract UCs. 2. We implement our method in TRP++, and we experimentally evaluate it. We make the source code of our solver available. Conceptually, under the frequently legitimate assumption that a system description can be translated into an LTL formula, our results extend to vacuity for LTL [BBDER01,KV03,AFF⁺03,FKSFV08,Kup06].

Related Work In [CRST07] Cimatti et al. perform extraction of UCs for PSL to accelerate a PSL satisfiability solver by performing Boolean abstraction. Their notion of UCs is coarser than ours and their solver is based on BDDs and on SAT. An investigation of notions of UCs for LTL including the relation between UCs and vacuity is performed in [Sch12b]. No implementation or experimental results are reported, and TR is not considered. Hantry et al. suggest a method to extract UCs for LTL in a tableau-based solver [HH11]. No implementation or experiments are reported. Awad et al. [AGH⁺12] use tableaux to extract UCs in the context of synthesizing business process templates. The description of the method is sketchy and incomplete, the notion of UC appears to be one of a subset of a set of formulas, and no detailed experimental evaluation is carried out. In [CMT11] Cimatti et al. show how to prove and explain unfeasibility of message sequence charts for networks of hybrid automata. They consider a different specification language and use an SMT-based algorithm. Some work deals with unrealizable rather than unsatisfiable cores. [CRST08] handles specifications in GR(1), which is a proper subset of LTL. Könighofer et al. present methods to help debugging unrealizable specifications by extracting unrealizable cores and simulating counterstrategies [KHB09] as well as performing error localization using model-based diagnosis [KHB10]. Raman and Kress-Gazit [RKG11] present a tool that points out unrealizable cores in the context of robot control. [Sch12b] explores more fine-grained notions of unrealizable cores than [CRST08,KHB09].

Structure of the Paper Section 2 starts with preliminaries. TR and its clausal normal form SNF are introduced in Sec. 3. In Sec. 4 we describe the construction of a resolution graph and its use to obtain a UC. The UCs obtained in Sec. 4 are lifted from SNF to LTL in Sec. 5. In Sec. 6 we provide examples that illustrate why these UCs are useful and how to obtain them. We discuss our implementation and experimental evaluation in Sec. 7. Section 8 concludes. Due to space constraints proofs are sketched or omitted. For a full version [Sch12a] of

this paper including proofs and for our implementation, examples, and log files see [pap].

2 Preliminaries

We use a standard version of LTL, see, e.g., [Eme90]. Let \mathbb{B} be the set of Booleans, and let AP be a finite set of atomic propositions. The set of *LTL formulas* is constructed inductively as follows. The Boolean constants 0 (false), 1 (true) $\in \mathbb{B}$ and any atomic proposition $p \in AP$ are LTL formulas. If ψ, ψ' are LTL formulas, so are $\neg\psi$ (not), $\psi \vee \psi'$ (or), $\psi \wedge \psi'$ (and), $\mathbf{X}\psi$ (next time), $\psi \mathbf{U} \psi'$ (until), $\psi \mathbf{R} \psi'$ (releases), $\mathbf{F}\psi$ (finally), and $\mathbf{G}\psi$ (globally). We use $\psi \rightarrow \psi'$ (implies) as an abbreviation for $\neg\psi \vee \psi'$.

3 Temporal Resolution (TR)

In this section we describe TR [FDP01] as implemented in TRP++ [HK03, HK04, trp]. We first explain the clausal normal form TR is based on. In Sec. 3.2 we provide a concise description of TR as required for the purposes of this paper. In Sec. 3.3 we give some intuition on how TR works with a slant towards BDD-based symbolic model checking (e.g., [BCM⁺92, CGH97, CGP01]). We wish to emphasize that TR is an existing technique that has been developed since the early 1990s [Fis91]. Our contribution in this paper is an extension of TR that allows to extract UCs from a run of TRP++ on an unsatisfiable LTL formula. Hence, while in this section we do our best to provide both a precise description of TR as well as some intuition on how it works, space constraints limit the extent of this description. Therefore, we refer readers interested in more details, a more extensive explanation, or correctness proofs of TR to [FDP01] for a general overview, to [Dix98, Dix97, Dix96, Dix95] for details on loop search, and to [HK03, HK04, trp] for the implementation in TRP++.

3.1 Separated Normal Form (SNF)

TR works on formulas in a clausal normal form called separated normal form (SNF) [Fis91, FN92, FDP01]. For any atomic proposition $p \in AP$ p and $\neg p$ are *literals*. Let $p_1, \dots, p_n, q_1, \dots, q_{n'}, l$ with $0 \leq n, n'$ be literals such that p_1, \dots, p_n and $q_1, \dots, q_{n'}$ are pairwise different. Then 1. $(p_1 \vee \dots \vee p_n)$ is an *initial clause*; 2. $(\mathbf{G}((p_1 \vee \dots \vee p_n) \vee (\mathbf{X}(q_1 \vee \dots \vee q_{n'}))))$ is a *global clause*; and 3. $(\mathbf{G}((p_1 \vee \dots \vee p_n) \vee (\mathbf{F}(l))))$ is an *eventuality clause*. l is called an *eventuality literal*. As usual an empty disjunction (resp. conjunction) stands for 0 (resp. 1). $()$ or $(\mathbf{G}())$, denoted \square , stand for 0 or $\mathbf{G}(0)$ and are called *empty clause*. The set of all SNF clauses is denoted \mathbb{C} . Let c_1, \dots, c_n with $0 \leq n$ be SNF clauses. Then $\bigwedge_{1 \leq i \leq n} c_i$ is an LTL formula in *SNF*. Every LTL formula ϕ can be transformed into an equisatisfiable formula ϕ' in SNF [FDP01].

rule	premise 1	pa- rt.	premise 2	pa- rt.	conclusion	pa- rt.	p.1 - c	p.2 - c	vt. c
saturation									
init-ii	$(P \vee l)$	M	$(\neg l \vee Q)$	M	$(P \vee Q)$	M	✓	✓	✓
init-in	$(P \vee l)$	M	$(\mathbf{G}(\neg l \vee Q))$	M	$(P \vee Q)$	M	✓	✓	✓
step-nn	$(\mathbf{G}(P \vee l))$	M	$(\mathbf{G}(\neg l \vee Q))$	M	$(\mathbf{G}(P \vee Q))$	M	✓	✓	✓
step-nx	$(\mathbf{G}(P \vee l))$	M	$(\mathbf{G}((Q) \vee (\mathbf{X}(\neg l \vee R))))$	M	$(\mathbf{G}((Q) \vee (\mathbf{X}(P \vee R))))$	M	✓	✓	✓
step-xx	$(\mathbf{G}(P) \vee (\mathbf{X}(Q \vee l)))$	ML	$(\mathbf{G}(R) \vee (\mathbf{X}(\neg l \vee S)))$	ML	$(\mathbf{G}((P \vee R) \vee (\mathbf{X}(Q \vee S))))$	ML	✓	✓	✓
augmentation									
aug1	$(\mathbf{G}((P) \vee (\mathbf{F}(l))))$			M	$(\mathbf{G}(P \vee l \vee wl))$	M	✓	—	✓
aug2	$(\mathbf{G}((P) \vee (\mathbf{F}(l))))$			M	$(\mathbf{G}((\neg wl) \vee (\mathbf{X}(l \vee wl))))$	M	✗	—	✓
BFS loop search									
BFS-loop- it-init-x	$c \equiv (\mathbf{G}((P) \vee (\mathbf{X}(Q))))$ with $ Q > 0$			M	c	L	✓	—	✓
BFS-loop- it-init-n	$(\mathbf{G}(P))$			M	$(\mathbf{G}((0) \vee (\mathbf{X}(P))))$	L	✓	—	✓
BFS-loop- it-init-c	$(\mathbf{G}(P))$	L'	$(\mathbf{G}((Q) \vee (\mathbf{F}(l))))$	M	$(\mathbf{G}((0) \vee (\mathbf{X}(P \vee l))))$	L	✗	✗	✓
BFS-loop- it-sub	$c \equiv (\mathbf{G}(P))$ with $c \rightarrow (\mathbf{G}(Q))$			L	$(\mathbf{G}((0) \vee (\mathbf{X}(Q \vee l))))$ generated by BFS-loop-it-init-c	L	✓	—	✗
BFS-loop- conclusion1	$(\mathbf{G}(P))$	L	$(\mathbf{G}((Q) \vee (\mathbf{F}(l))))$	M	$(\mathbf{G}(P \vee Q \vee l))$	M	✓	✓	✓
BFS-loop- conclusion2	$(\mathbf{G}(P))$	L	$(\mathbf{G}((Q) \vee (\mathbf{F}(l))))$	M	$(\mathbf{G}((\neg wl) \vee (\mathbf{X}(P \vee l))))$	M	✓	✗	✓

Table 1. Production rules used in TRP++. Let $P \equiv p_1 \vee \dots \vee p_n$, $Q \equiv q_1 \vee \dots \vee q_{n'}$, $R \equiv r_1 \vee \dots \vee r_{n''}$, and $S \equiv s_1 \vee \dots \vee s_{n'''}$.

3.2 TR in TRP++

The production rules of TRP++ are shown in Tab. 1. The 1st column assigns a name to a production rule. The 2nd and 4th columns list the premises. The 6th column gives the conclusion. Columns 3, 5, and 7 are described below. Columns 8–10 become relevant only in later sections.

Algorithm 1 provides a high level view of TR in TRP++ [HK04]. The algorithm takes a set of starting clauses C in SNF as input. It returns *unsat* if C is found to be unsatisfiable (by deriving \square) and *sat* otherwise. Resolution between two initial or two global clauses or between an initial and a global clause is performed by a straightforward extension of propositional resolution (e.g., [Rob65, FM09, BG01]). The corresponding production rules are listed under *saturation* in Tab. 1. Given a set of SNF clauses C we say that one *saturates* C if one applies these production rules to clauses in C until no new clauses are generated. Resolution between a set of initial and global clauses and an eventuality clause with eventuality literal l requires finding a set of global clauses that allows to infer conditions under which $\mathbf{XG}\neg l$ holds. Such a set of clauses is called a *loop* in $\neg l$. Loop search involves all production rules in Tab. 1 except `init-ii`, `init-in`, `step-nn`, and `step-nx`.

In line 1 Alg. 1 initializes M with the set of starting clauses and terminates iff one of these is the empty clause. Then, in line 2, it saturates M (terminating iff the empty clause is generated). In line 3 it *augments* M by applying production rule `aug1` to each eventuality clause in M and `aug2` once per eventuality literal in M , where wl is a fresh proposition. This is followed by another round of saturation in line 4. From now on Alg. 1 alternates between searching for a loop for some eventuality clause c (lines 9–18) and saturating M if loop search has

Algorithm 1: LTL satisfiability checking via TR in TRP++.

Input: A set of SNF clauses C . **Output:** *Unsat* if C is unsatisfiable; *sat* otherwise.

```
1  $M \leftarrow C$ ; if  $\Box \in M$  then return unsat;
2 saturate( $M$ ); if  $\Box \in M$  then return unsat;
3 augment( $M$ );
4 saturate( $M$ ); if  $\Box \in M$  then return unsat;
5  $M' \leftarrow \emptyset$ ;
6 while  $M' \neq M$  do
7    $M' \leftarrow M$ ;
8   for  $c \in C$  .  $c$  is an eventuality clause do
9      $C' \leftarrow \{\Box\}$ ;
10    repeat
11      initialize-BFS-loop-search-iteration( $M, c, C', L$ );
12      saturate-step-xx( $L$ );
13       $C' \leftarrow \{c' \in L \mid c' \text{ has empty } \mathbf{X} \text{ part}\}$ ;
14       $C'' \leftarrow \{(\mathbf{G}(Q)) \mid (\mathbf{G}((0) \vee (\mathbf{X}(Q \vee l)))) \in L \text{ generated by } \boxed{\text{BFS-loop-it-init-c}}\}$ ;
15       $\text{found} \leftarrow \text{subsumes}(C', C'')$ ;
16    until found or  $C' = \emptyset$ ;
17    if found then
18      derive-BFS-loop-search-conclusions( $c, C', M$ );
19      saturate( $M$ ); if  $\Box \in M$  then return unsat;
20 return sat;
```

generated new clauses (line 19). It terminates, if either the empty clause was derived (line 19) or if no new clauses were generated (line 20).

Loop search for some eventuality clause c may take several *iterations* (lines 11–15). Each loop search iteration uses saturation restricted to $\boxed{\text{step-xx}}$ as a subroutine (line 12). Therefore, each loop search iteration has its own set of clauses L in which it works. We call M and L *partitions*. Columns 3, 5, and 7 in Tab. 1 indicate whether a premise (resp. conclusion) of a production rule is taken from (resp. put into) the main partition (M), the loop partition of the current loop search iteration (L), the loop partition of the previous loop search iteration (L'), or either of M or L as long as premises and conclusion are in the same partition (ML). In line 11 partition L of a loop search iteration is initialized by applying production rule $\boxed{\text{BFS-loop-it-init-x}}$ once to each global clause with non-empty \mathbf{X} part in M , rule $\boxed{\text{BFS-loop-it-init-n}}$ once to each global clause with empty \mathbf{X} part in M , and rule $\boxed{\text{BFS-loop-it-init-c}}$ once to each global clause with empty \mathbf{X} part in the partition of the previous loop search iteration L' . Notice that by construction at this point L contains only global clauses with non-empty \mathbf{X} part. Then L is saturated using only rule $\boxed{\text{step-xx}}$ (line 12). A loop has been found iff each global clause with empty \mathbf{X} part that was derived in the previous loop search iteration is subsumed by at least one global clause with empty \mathbf{X} part that was derived in the current loop search iteration (lines 13–15). Subsumption between a pair of clauses corresponds to an instance of production rule $\boxed{\text{BFS-loop-it-sub}}$; note, though, that this rule does not produce a new clause but records a relation between two clauses to be used later for extraction of a UC. Loop search for c terminates, if either a loop has been found or no clauses with empty \mathbf{X} part were derived (line 16). If a loop has been found,

rules $\boxed{\text{BFS-loop-conclusion1}}$ and $\boxed{\text{BFS-loop-conclusion2}}$ are applied once to each global clause with empty **X** part that was derived in the current loop search iteration (line 18) to obtain the loop search conclusions for the main partition.

3.3 TR — Some Intuition

The following explanation is partly based on the correctness proofs used in TR (e.g., [FDP01, Dix95]). At various points we draw parallels to BDD-based symbolic model checking (below shortened to “model checking”; e.g., [BCM⁺92, CGH97, CGP01]).

Transition Systems Given a set of atomic propositions AP a transition system $G = (V, E, I)$ is a directed graph with a finite set of vertices $V \subseteq 2^{AP}$, a set of directed edges $E \subseteq V \times V$, and a set of initial vertices $I \subseteq V$. A set of SNF clauses C induces a transition system as follows. C is partitioned into the 3 sets of initial clauses, global clauses with empty **X** part, and global clauses with non-empty **X** part (eventuality clauses are ignored). The set of vertices V is given by those valuations of AP that fulfill the bodies of the global clauses with empty **X** part. The set of edges is given by those pairs of vertices that fulfill the bodies of the global clauses with non-empty **X** part. The set of initial vertices is the subset of vertices that fulfill the initial clauses. Note that not all vertices may be reachable from an initial vertex. Given a set of SNF clauses C it is easy to see that its induced transition system contains an initialized infinite path that fulfills the eventuality clauses in C if and only if C is satisfiable.

Saturation Saturation adds clauses to a set of SNF clauses C such that the resulting induced transition system G' is restricted to those vertices of G that are the start of an infinite path in G . Resolution within each of the sets of initial clauses ($\boxed{\text{init-ii}}$), global clauses with empty **X** part ($\boxed{\text{step-nn}}$), and global clauses with non-empty **X** part ($\boxed{\text{step-xx}}$) — provided the result has a non-empty **X** part — induce the same transition system before and after such resolution. The same is true for resolution between a global clause with empty **X** part and a global clause with non-empty **X** part ($\boxed{\text{step-nx}}$) that results in a global clause with non-empty **X** part as well as for resolution between an initial clause and a global clause with empty **X** part ($\boxed{\text{init-in}}$). Note that while such resolution does not lead to a change in the induced transition system, these resolutions serve 2 purposes: 1. to generate clauses that are needed as input for resolution that does lead to a change in the induced transition system and 2. to generate the empty clause as a sign for unsatisfiability. Resolution between two global clauses with non-empty **X** part ($\boxed{\text{step-xx}}$) that results in a global clause with empty **X** part c may induce different transition systems before and after such resolution: before resolution is carried out, the induced transition system may have vertices not fulfilling the body of c , albeit with no outgoing edges; after resolution is carried out, the induced transition system will have no such vertices. This is also true for resolution between a global clause with empty **X** part and a global clause with non-empty **X** part ($\boxed{\text{step-nx}}$) that results in a global clause with empty

X part. As propositional resolution is complete and the above resolution steps are carried out until no new clauses are generated, at that point the induced transition system contains only vertices that are the start of an infinite path as stated above. This is somewhat comparable to model checking pruning a transition system to the set of vertices that are the start of an infinite path by repeatedly computing a backward image for a set of vertices and intersecting with that set of vertices until a fixed point is reached; note though, that model checking tends to proceed in a breadth-first manner, while TR is more flexible.

BFS Loop Search Given a set of SNF clauses C with an eventuality clause $(\mathbf{G}((P) \vee (\mathbf{F}(l))))$ BFS loop search adds clauses to C that remove those vertices from the induced transition system that do not fulfill P but cannot reach a vertex that fulfills l . Assume an eventuality clause $(\mathbf{G}((P) \vee (\mathbf{F}(l))))$. The first iteration of a BFS loop search for a loop in $\neg l$ produces global clauses with empty **X** part whose body is fulfilled by those vertices in the induced transition system G that can reach a vertex in G that fulfills l in one step. The second iteration extends that to one or two steps. This continues until a fixed point is reached where the bodies of the produced global clauses with empty **X** part are fulfilled by those vertices in the induced transition system G that can reach a vertex in G that fulfills l in one or more steps. The loop search conclusions then combine this new information with P and wl . In model checking the computation of vertices that can reach l in one or more steps by repeated computation of backward images is very similar; note, though, that in typical model checking with Büchi fairness (e.g., [BCM⁺92, CGH97, CGP01]) P and wl are not present.

High Level View In App. A we turn the above discussion into a high level view of TR in TRP++ and we discuss the relation to cycle detection algorithms in model checking.

4 UC Extraction

In this section we describe, given an unsatisfiable set of SNF clauses C , how to obtain a subset of C , C^{uc} , that is by itself unsatisfiable from an execution of Alg. 1. The general idea of the construction is unsurprising in that during the execution of Alg. 1 a resolution graph is built that records which clauses were used to generate other clauses (Def. 1). Then the resolution graph is traversed backwards from the empty clause to find the subset of C that was actually used to prove unsatisfiability (Def. 2). The main concern of Def. 1, 2, and their proof of correctness in Thm. 1 is therefore that/why certain parts of the TR proof do not need to be taken into account when determining C^{uc} .

Definition 1 (Resolution Graph). A resolution graph G is a directed graph consisting of 1. a set of vertices V , 2. a set of directed edges $E \subseteq V \times V$, 3. a labeling of vertices with SNF clauses $L_V : V \rightarrow \mathbb{C}$, and 4. a partitioning \mathcal{Q}^V of the set of vertices V into one main partition M^V and one partition L_i^V for each

BFS loop search iteration: $\mathcal{Q}^V : V = M^V \uplus L_0^V \uplus \dots \uplus L_n^V$. Let C be a set of SNF clauses. During an execution of Alg. 1 with input C a resolution graph G is constructed as follows.

In line 1 G is initialized: 1. V contains one vertex v per clause c in C : $V = \{v_c \mid c \in C\}$, 2. E is empty: $E = \emptyset$, 3. each vertex is labeled with the corresponding clause: $L_V : V \rightarrow C, L_V(v_c) = c$, and 4. the partitioning \mathcal{Q}^V contains only the main partition M^V , which contains all vertices: $\mathcal{Q}^V : M^V = V$.

Whenever a new BFS loop search iteration is entered (line 11), a new partition L_i^V is created and added to \mathcal{Q}^V . For each application of a production rule from Tab. 1 that either generates a new clause in partition M or L or is the first application of rule `BFS-loop-it-sub` to clause c'' in C'' in line 15: 1. if column 10 (vt. c) of Tab. 1 contains \checkmark , then a new vertex v is created for the conclusion c (which is a new clause), labeled with c , and put into partition M^V or L_i^V ; 2. if column 8 (p.1 - c) (resp. column 9 (p.2 - c)) contains \checkmark , then an edge is created from the vertex labeled with premise 1 (resp. premise 2) in partition M^V or L_i^V to the vertex labeled with the conclusion in partition M^V or L_i^V .

Definition 2 (UC in SNF). Let C be a set of SNF clauses to which Alg. 1 has been applied and shown unsatisfiability, let G be the resolution graph, and let v_\square be the (unique) vertex in the main partition M^V of the resolution graph G labeled with the empty clause \square . Let G' be the smallest subgraph of G that contains v_\square and all vertices in G (and the corresponding edges) that are backward reachable from v_\square . The UC of C in SNF, C^{uc} , is the subset of C such that there exists a vertex v in the subgraph G' , labeled with $c \in C$, and contained in the main partition M^V of G : $C^{uc} = \{c \in C \mid \exists v \in V_{G'} . L_V(v) = c \wedge v \in M^V\}$.

Theorem 1 (Unsatisfiability of UC in SNF). Let C be a set of SNF clauses to which Alg. 1 has been applied and shown unsatisfiability, and let C^{uc} be the UC of C in SNF. Then C^{uc} is unsat.

Assume for a moment that in columns 8 (p.1 - c) and 9 (p.2 - c) of Tab. 1 all \times are replaced with \checkmark , i.e., that each conclusion in the resolution graph is connected by an edge to each of its premises rather than only to a subset of them. In that case the UC in SNF according to Def. 2 would contain all clauses of the set of starting clauses C that contributed to deriving the empty clause and, hence, to establishing unsatisfiability of C . In that case it would follow directly from the correctness of TR that C^{uc} is unsatisfiable. In the proof (see App. B) it remains to show that not including an edge 1. from premise 1 to the conclusion for rule `aug2`, 2. from premise 2 to the conclusion for rule `BFS-loop-conclusion2`, 3. from premise 2 to the conclusion for rule `BFS-loop-it-init-c`, and 4. from premise 1 to the conclusion for rule `BFS-loop-it-init-c` in the resolution graph G maintains the fact that the resulting C^{uc} is unsatisfiable.

By taking the fact that each vertex in the resolution graph has at most 2 incoming edges into account, the first part of the following Prop. 1 is immediate from Def. 1 and 2. The second part is obtained by bounding the number of 1. different clauses in each partition, 2. iterations in each loop search by the

Subf.	Prop.	SNF Clauses (+ polarity occurrences)	SNF Clauses (− polarity occurrences)
$1/0/p$	$1/0/p$	—	—
$\neg\psi$	$x_{\neg\psi}$	$(\mathbf{G}(x_{\neg\psi} \rightarrow (\neg \boxed{x_{\psi}})))$	$(\mathbf{G}((\neg x_{\neg\psi}) \rightarrow \boxed{x_{\psi}}))$
$\psi \wedge \psi'$	$x_{\psi \wedge \psi'}$	$(\mathbf{G}(x_{\psi \wedge \psi'} \rightarrow \boxed{x_{\psi}})), (\mathbf{G}(x_{\psi \wedge \psi'} \rightarrow \boxed{x_{\psi'}}))$	$(\mathbf{G}((\neg x_{\psi \wedge \psi'}) \rightarrow ((\neg \boxed{x_{\psi}}) \vee (\neg \boxed{x_{\psi'}}))))$
$\psi \vee \psi'$	$x_{\psi \vee \psi'}$	$(\mathbf{G}(x_{\psi \vee \psi'} \rightarrow (\boxed{x_{\psi}} \vee \boxed{x_{\psi'}})))$	$(\mathbf{G}((\neg x_{\psi \vee \psi'}) \rightarrow (\neg \boxed{x_{\psi}}))), (\mathbf{G}((\neg x_{\psi \vee \psi'}) \rightarrow (\neg \boxed{x_{\psi'}})))$
$\mathbf{X}\psi$	$\mathbf{X}x_{\psi}$	$(\mathbf{G}(\mathbf{X}x_{\psi} \rightarrow (\mathbf{X}\boxed{x_{\psi}})))$	$(\mathbf{G}((\neg \mathbf{X}x_{\psi}) \rightarrow (\mathbf{X}\neg \boxed{x_{\psi}})))$
$\mathbf{G}\psi$	$\mathbf{G}x_{\psi}$	$(\mathbf{G}(x_{\mathbf{G}\psi} \rightarrow (\mathbf{X}x_{\mathbf{G}\psi}))), (\mathbf{G}(x_{\mathbf{G}\psi} \rightarrow \boxed{x_{\psi}}))$	$(\mathbf{G}((\neg x_{\mathbf{G}\psi}) \rightarrow (\mathbf{F}\neg \boxed{x_{\psi}})))$
$\mathbf{F}\psi$	$\mathbf{F}x_{\psi}$	$(\mathbf{G}(x_{\mathbf{F}\psi} \rightarrow (\mathbf{F}\boxed{x_{\psi}})))$	$(\mathbf{G}((\neg x_{\mathbf{F}\psi}) \rightarrow (\mathbf{X}\neg \boxed{x_{\psi}}))), (\mathbf{G}((\neg x_{\mathbf{F}\psi}) \rightarrow (\neg \boxed{x_{\psi}})))$
$\psi \mathbf{U} \psi'$	$x_{\psi \mathbf{U} \psi'}$	$(\mathbf{G}(x_{\psi \mathbf{U} \psi'} \rightarrow (\boxed{x_{\psi'}} \vee \boxed{x_{\psi}}))),$ $(\mathbf{G}(x_{\psi \mathbf{U} \psi'} \rightarrow (\boxed{x_{\psi'}} \vee (\mathbf{X}x_{\psi \mathbf{U} \psi'}))),$ $(\mathbf{G}(x_{\psi \mathbf{U} \psi'} \rightarrow (\mathbf{F}x_{\psi'})))$	$(\mathbf{G}((\neg x_{\psi \mathbf{U} \psi'}) \rightarrow (\neg \boxed{x_{\psi'}}))),$ $(\mathbf{G}((\neg x_{\psi \mathbf{U} \psi'}) \rightarrow ((\neg \boxed{x_{\psi'}}) \vee (\mathbf{X}\neg x_{\psi \mathbf{U} \psi'}))))$
$\psi \mathbf{R} \psi'$	$x_{\psi \mathbf{R} \psi'}$	$(\mathbf{G}(x_{\psi \mathbf{R} \psi'} \rightarrow \boxed{x_{\psi'}})),$ $(\mathbf{G}(x_{\psi \mathbf{R} \psi'} \rightarrow (\boxed{x_{\psi}} \vee (\mathbf{X}x_{\psi \mathbf{R} \psi'}))))$	$(\mathbf{G}((\neg x_{\psi \mathbf{R} \psi'}) \rightarrow ((\neg \boxed{x_{\psi'}}) \vee (\neg \boxed{x_{\psi}}))),$ $(\mathbf{G}((\neg x_{\psi \mathbf{R} \psi'}) \rightarrow ((\neg \boxed{x_{\psi'}}) \vee (\mathbf{X}\neg x_{\psi \mathbf{R} \psi'}))),$ $(\mathbf{G}((\neg x_{\psi \mathbf{R} \psi'}) \rightarrow (\mathbf{F}\neg \boxed{x_{\psi'}})))$

Table 2. Translation from LTL to SNF.

length of the longest monotonically increasing sequence of Boolean formulas over AP , and 3. loop searches by the number of different loop search conclusions.

Proposition 1 (Complexity of UC Extraction). *Let C be a set of SNF clauses to which Alg. 1 is applied and shows unsatisfiability. Construction and backward traversal of the resolution graph and, hence, construction of C^{uc} according to Def. 2 can be performed in time $\mathcal{O}(|V|)$ in addition to the time required to run Alg. 1. $|V|$ is at most exponential in $|AP| + \log(|C|)$.*

5 From LTL to SNF and Back

We use a structure-preserving translation to translate an LTL formula into a set of SNF clauses, which slightly differs from the translation suggested in [FDP01]. It is well known that ϕ and $SNF(\phi)$ according to Def. 3 are equisatisfiable and that a satisfying assignment for ϕ (resp. $SNF(\phi)$) can be extended (resp. restricted) to a satisfying assignment of $SNF(\phi)$ (resp. ϕ).

Definition 3 (Translation from LTL to SNF). *Let ϕ be an LTL formula over atomic propositions AP , and let $X = \{x, x', \dots\}$ be a set of fresh atomic propositions not in AP . Assign each occurrence of a subformula ψ in ϕ a Boolean value or a proposition according to col. 2 of Tab. 2, which is used to reference ψ in the SNF clauses for its superformula. Moreover, assign each occurrence of ψ a set of SNF clauses according to col. 3 or 4 of Tab. 2. Let $SNF_{aux}(\phi)$ be the set of all SNF clauses obtained from ϕ that way. Then the SNF of ϕ is defined as $SNF(\phi) \equiv x_{\phi} \wedge \bigwedge_{c \in SNF_{aux}(\phi)} c$.*

In the following Def. 4 we describe how to map a UC in SNF back to a UC in LTL. The main idea in its proof of correctness (Thm. 2) is to compare the SNF of ϕ and ϕ^{uc} by partitioning the SNF clauses into three sets: one that is shared by the two SNFs, one that replaces some occurrences of propositions in $SNF(\phi)$ with 1 or 0, and one whose clauses are only in $SNF(\phi)$. Then one can show that the UC of ϕ in SNF must be contained in the first partition.

Definition 4 (Mapping a UC in SNF to a UC in LTL). Let ϕ be an unsatisfiable LTL formula, let $\text{SNF}(\phi)$ be its SNF, and let C^{uc} be the UC of $\text{SNF}(\phi)$ in SNF. Then the UC of ϕ in LTL, ϕ^{uc} , is obtained as follows. For each positive (resp. negative) polarity occurrence of a proper subformula ψ of ϕ with proposition x_ψ according to Tab. 2, replace ψ in ϕ with 1 (resp. 0) iff C^{uc} contains no clause with an occurrence of proposition x_ψ that is marked [blue boxed](#) in Tab. 2. (We are sloppy in that we “replace” subformulas of replaced subformulas, while in effect they simply vanish.)

Theorem 2 (Unsatisfiability of UC in LTL). Let ϕ be an unsatisfiable LTL formula, and let ϕ^{uc} be the UC of ϕ in LTL. Then ϕ^{uc} is unsat.

Remark 1. In Def. 10 of [Sch12b] a UC of an unsatisfiable formula in LTL is obtained by replacing some occurrences of positive polarity subformulas with 1 and some occurrences of negative polarity subformulas with 0 while maintaining unsatisfiability. By construction in Def. 4 and with Thm. 2 it is immediate that a UC in LTL according to Def. 4 above is a UC according to Def. 10 of [Sch12b].

6 Examples

In this section we first present examples of using UCs for LTL to help understanding a specification given in LTL. Then we show an example of TR with the corresponding resolution graph and UC extraction in SNF.

Using UCs in LTL to Help Understanding LTL Specifications We start with a toy example and then proceed to a more realistic one. Except for minor rewriting, all UCs in this section were obtained with our implementation. The first example (1a)–(1c) is based on [JB06]. We would like to see whether a *req* (request) can be issued (1d). This is impossible, as (1a) requires a *req* to be followed by 3 *gnts* (grant), whereas (1b) forbids two subsequent *gnts*. The UC in (2) clearly shows this.

$$\begin{array}{ll}
(\mathbf{G}(req \rightarrow ((\mathbf{X}gnt) \wedge (\mathbf{XX}gnt) \wedge (\mathbf{XXX}gnt)))) & (1a) \\
\wedge (\mathbf{G}(gnt \rightarrow \mathbf{X}\neg gnt)) & (1b) \\
\wedge (\mathbf{G}(cancel \rightarrow \mathbf{X}((\neg gnt)\mathbf{U}go))) & (1c) \\
\wedge (\mathbf{F}req) & (1d)
\end{array}
\qquad
\begin{array}{ll}
(\mathbf{G}(req \rightarrow ((\mathbf{X}gnt) \wedge (\mathbf{XX}gnt)))) & (2a) \\
\wedge (\mathbf{G}(gnt \rightarrow \mathbf{X}\neg gnt)) & (2b) \\
\wedge (\mathbf{F}req) & (2c)
\end{array}$$

The 2nd example (3) is adapted from a lift specification in [Har05] (we used a somewhat similar example in [Sch12b]). The lift has two floors, indicated by f_0 and f_1 . On each floor there is a button to call the lift (b_0 , b_1). *sb* is 1 if some button is pressed. If the lift moves up, then *up* must be 1; if it moves down, then *up* must be 0. *u* switches turns between actions by users of the lift (*u* is 1) and

actions by the lift (u is 0). For a more detailed explanation we refer to [Har05].

$$\begin{aligned}
& (\neg u) \wedge (f_0) \wedge (\neg b_0) \wedge (\neg b_1) \wedge (\neg up) & (3a) \\
& \wedge (\mathbf{G}((u \rightarrow \neg \mathbf{X}u) \wedge ((\neg \mathbf{X}u) \rightarrow u))) & (3b) \\
& \wedge (\mathbf{G}(f_0 \rightarrow \neg f_1)) & (3c) \\
& \wedge (\mathbf{G}((f_0 \rightarrow \mathbf{X}(f_0 \vee f_1)) \wedge (f_1 \rightarrow \mathbf{X}(f_0 \vee f_1)))) & (3d) \\
& \wedge (\mathbf{G}(u \rightarrow ((f_0 \rightarrow \mathbf{X}f_0) \wedge ((\mathbf{X}f_0) \rightarrow f_0) \wedge (f_1 \rightarrow \mathbf{X}f_1) \wedge ((\mathbf{X}f_1) \rightarrow f_1)))) & (3e) \\
& \wedge (\mathbf{G}(((\neg u) \rightarrow ((b_0 \rightarrow \mathbf{X}b_0) \wedge ((\mathbf{X}b_0) \rightarrow b_0) \wedge (b_1 \rightarrow \mathbf{X}b_1) \wedge ((\mathbf{X}b_1) \rightarrow b_1)))))) & (3f) \\
& \wedge (\mathbf{G}(((b_0 \wedge \neg f_0) \rightarrow \mathbf{X}b_0) \wedge ((b_1 \wedge \neg f_1) \rightarrow \mathbf{X}b_1))) & (3g) \\
& \wedge (\mathbf{G}((f_0 \wedge \mathbf{X}f_0) \rightarrow ((up \rightarrow \mathbf{X}up) \wedge ((\mathbf{X}up) \rightarrow up)))) & (3h) \\
& \wedge (\mathbf{G}((f_1 \wedge \mathbf{X}f_1) \rightarrow ((up \rightarrow \mathbf{X}up) \wedge ((\mathbf{X}up) \rightarrow up)))) & (3i) \\
& \wedge (\mathbf{G}(((f_0 \wedge \mathbf{X}f_1) \rightarrow up) \wedge ((f_1 \wedge \mathbf{X}f_0) \rightarrow \neg up))) & (3j) \\
& \wedge (\mathbf{G}((sb \rightarrow (b_0 \vee b_1)) \wedge ((b_0 \vee b_1) \rightarrow sb))) & (3k) \\
& \wedge (\mathbf{G}(((f_0 \wedge \neg sb) \rightarrow (f_0 \mathbf{U}(sb \mathbf{R}((\mathbf{F}f_0) \wedge (\neg up)))))) & (3l) \\
& \wedge (\mathbf{G}(((f_1 \wedge \neg sb) \rightarrow (f_1 \mathbf{U}(sb \mathbf{R}((\mathbf{F}f_0) \wedge (\neg up)))))) & (3m) \\
& \wedge (\mathbf{G}((b_0 \rightarrow \mathbf{F}f_0) \wedge (b_1 \rightarrow \mathbf{F}f_1))) & (3n)
\end{aligned}$$

We first assume that an engineer is interested in seeing whether it is possible that b_1 is continuously pressed (4). As the UC (5) shows, this is impossible as b_1 must be 0 at the beginning.

$$\mathbf{G}b_1 \quad (4) \quad (\neg b_1) \wedge \mathbf{G}b_1 \quad (5)$$

Now the engineer modifies her query such that b_1 is pressed only from time point 1 on (6). As shown by the UC in (7) that turns out to be impossible, too.

$$\mathbf{XG}b_1 \quad (6) \quad (\neg u) \wedge ((\neg b_1) \wedge ((\mathbf{G}((\neg u) \rightarrow ((\mathbf{X}b_1) \rightarrow b_1))) \wedge (\mathbf{XG}b_1))) \quad (7)$$

The engineer now tries to have b_1 pressed only from time point 2 on and, again, obtains a UC. She becomes suspicious and checks whether b_1 can be pressed at all (8). She now sees that b_1 cannot be pressed at all and, therefore, this specification of a lift must contain a bug. She can now use the UC in (9a)–(9f) to track down the problem. This example illustrates the use of UCs for debugging, as (9a)–(9f) is significantly smaller than (3).

$$\begin{array}{llll}
& (f_0) \wedge (\neg b_1) \wedge (\neg up) & (9a) & \\
\mathbf{F}b_1 \quad (8) & \wedge (\mathbf{G}(f_0 \rightarrow \neg f_1)) & (9b) & \wedge (\mathbf{G}((f_0 \wedge \mathbf{X}f_1) \rightarrow up)) \quad (9e) \\
& \wedge (\mathbf{G}(f_0 \rightarrow \mathbf{X}(f_0 \vee f_1))) & (9c) & \wedge (\mathbf{G}(b_1 \rightarrow \mathbf{F}f_1)) \quad (9f) \\
& \wedge (\mathbf{G}((f_0 \wedge \mathbf{X}f_0) \rightarrow ((\mathbf{X}up) \rightarrow up))) & (9d) & \wedge (\mathbf{F}(b_1)) \quad (9g)
\end{array}$$

TR, Resolution Graph, and UC Extraction In Fig. 1 we show an example of an execution of the TR algorithm with the corresponding resolution graph and UC extraction in SNF. The set of starting clauses C to be solved is $\mathbf{G}(a \vee \neg b)$, $\mathbf{G}(a \vee b \vee \mathbf{X}(a \vee b))$, $\mathbf{G}((\neg a) \vee \mathbf{X}a)$, $\mathbf{G}((\neg a) \vee \mathbf{F}\neg a)$, shown in the first row from the bottom in the rectangle shaded in light red. In Fig. 1 TR generally proceeds from bottom to top; in the top right corner the empty clause \square is generated, indicating unsatisfiability. Clauses are connected with directed edges from premises to conclusions according to columns 8, 9 in Tab. 1. Edges are labeled with production rules, where “BFS-loop” is abbreviated to “loop”, “init” to “i”, and “conclusion” to “conc”. Saturation in line 2 of Alg. 1 produces $\mathbf{G}(a \vee b \vee \mathbf{X}a)$ in the 2nd row from the bottom.¹ The other 2 clauses in that row are generated by

¹ While it may seem that some clauses are not considered for loop initialization or saturation, this is due to either subsumption of one clause by another (e.g.,

augmentation (line 3). The following saturation (line 4) produces no new clauses. The dark green shaded rectangle is the loop partition for the first loop search iteration. Row 3 contains the clauses obtained by initialization of the BFS loop search iteration (line 11). Row 4 then contains the clauses generated from those in row 3 by saturation restricted to step-xx (line 12). The subsumption test fails in this iteration, as none of the clauses in row 4 subsumes the empty clause (lines 13–15). The light green shaded rectangle is the loop partition for the next loop search iteration. Row 5 contains the clauses obtained by initialization and row 6 those obtained from them by restricted saturation. This time the subsumption test succeeds, and the loop search conclusions are shown in row 7 (line 18). Finally, row 8 contains the derivation of the empty clause \square via saturation (line 19). The thick, dotted, blue clauses and edges show the part of the resolution graph that is backward reachable from \square . As all starting clauses in C are backward reachable from \square , the UC of C in SNF is C (note that this example serves to illustrate the mechanism rather than the benefit of UC extraction).

7 Experimental Evaluation

Our implementation, examples, and log files are available from [\[pap\]](#).

Implementation In a recent experimental evaluation of solvers for satisfiability of propositional LTL [SD11] TRP++ proved to be competitive. It is available as source code [\[trp\]](#). We therefore chose TRP++ as the basis for our implementation. TRP++ provides a translation from LTL to SNF via an external tool. To facilitate tracing a UC in SNF back to the input formula in LTL we implemented a translator from LTL to SNF inside TRP++, which reimplements ideas from the external translator. We used parts of TSPASS [LH10] for our implementation. For data structures we used C++ STL containers, for graph operations the Boost Graph Library [\[boo\]](#).

Benchmarks Our examples are based on [SD11]. In categories **crafted** and **random** and in family **forobots** we considered all unsatisfiable instances from [SD11]. The version of **alaska_lift** used here contains a small bug fix: in [WDMR08, SD11] the subformula $\mathbf{X}u$ was erroneously written as literal Xu . Combining 2 variants of **alaska_lift** with 3 different scenarios we obtain 6 subfamilies of **alaska_lift**. For **anzu_genbuf** we invented 3 scenarios to obtain 3 subfamilies. For all benchmark families that consist of a sequence of instances of increasing difficulty we stopped after two instances that could not be solved due to time or memory out. Some instances were simplified to 0 during the translation from LTL to SNF; these instances were discarded. In Tab. 3 we give an overview of the benchmark families. Columns 1–3 give the category, name, and the source

$\mathbf{G}(a \vee b \vee \mathbf{X}(a \vee b))$ by $\mathbf{G}(a \vee b \vee \mathbf{X}a)$ or the fact that TRP++ uses *ordered* resolution (e.g., $\mathbf{G}(a \vee b \vee \mathbf{X}a)$ with $\mathbf{G}(\neg w\bar{a} \vee \mathbf{X}((\neg a) \vee w\bar{a}))$; [HK03, BG01]). Both are issues of completeness of TR and, therefore, not discussed in this paper.

family	source	# solved no UC UC	largest solved
application			
alaska_lift	[Har05, WDMR08]	73 71	4605
anzu_genbuf	[BGJ ⁺ 07]	16 16	2676
forobots	[BDF09]	25 25	635
crafted			
schuppan_O1formula	[SD11]	21 21	1606
schuppan_O2formula	[SD11]	8 8	91
schuppan_phltl	[SD11]	4 4	125
random			
rozier_random	[RV10]	66 66	157
trp	[HS02]	397 397	1421

Table 3. Overview of benchmark families.




Fig. 2. (a) and (b): overhead incurred by UC extraction in terms of run time (in seconds) and memory (in GB) with no UC extraction on the x-axis and UC extraction on the y-axis. The off-center diagonal shows where $y = 2x$. (c): size reduction obtained by UC extraction. The x-axis shows the sizes of the input formulas, the y-axis shows the sizes of the UCs. Size is measured as the number of nodes in the syntax trees.

Discussion Our data show that extraction of UCs is possible with quite acceptable overhead in run time and memory usage (Fig. 2 (a), (b)). In particular, out of the 746 instances we considered with UC extraction disabled, 44 were simplified to 0 in the translation to SNF, 610 were shown to be unsatisfiable by TR, and 92 remained unsolved. Enabling UC extraction results in 2 time or memory outs out of 610 instances. The resulting UCs are often significantly smaller than the input formula (Fig. 2 (c)).

8 Conclusions

In this paper we showed how to obtain UCs for LTL via temporal resolution, and we demonstrated with an implementation in **TRP++** that UC extraction can be performed efficiently. The resulting UCs are significantly smaller than the corresponding input formulas. The similarity of temporal resolution and BDD-based algorithms at a high level (Sec. 3, App. A) and work on resolution with BDDs ([JSB06]) suggests to explore whether computation of UCs is feasible for BDD-based algorithms. Another direction for transfer of our results is resolution-based computation of unrealizable cores [Noë95]. An immediate possibility to optimize the UCs we obtain is minimization by repeating extraction of UCs until a fixed point is reached and subsequent attempts to delete remaining clauses [ZM03b].

Acknowledgements I thank B. Konev and M. Ludwig for making TRP++ and TSPASS including their LTL translators available. I also thank A. Cimatti for bringing up the subject of temporal resolution. Initial parts of the work were performed while working under a grant by the Provincia Autonoma di Trento (project EMTELOS).

References

- AFF⁺03. R. Armoni, L. Fix, A. Flaisher, O. Grumberg, N. Piterman, A. Tiemeyer, and M. Vardi. Enhanced vacuity detection in linear temporal logic. In W. Hunt Jr. and F. Somenzi, editors, *CAV*, volume 2725 of *Lecture Notes in Computer Science*, pages 368–380. Springer, 2003.
- AGH⁺12. A. Awad, R. Goré, Z. Hou, J. Thomson, and M. Weidlich. An iterative approach to synthesize business process templates from compliance rules. *Inf. Syst.*, 37(8):714–736, 2012.
- BBDER01. I. Beer, S. Ben-David, C. Eisner, and Y. Rodeh. Efficient detection of vacuity in temporal model checking. *Formal Methods in System Design*, 18(2):141–163, 2001.
- BCM⁺92. J. Burch, E. Clarke, K. McMillan, D. Dill, and L. Hwang. Symbolic model checking: 10^{20} states and beyond. *Inf. Comput.*, 98(2):142–170, 1992.
- BDF09. A. Behdenna, C. Dixon, and M. Fisher. Deductive verification of simple foraging robotic behaviours. *International Journal of Intelligent Computing and Cybernetics*, 2(4):604–643, 2009.
- BDTW93. R. Bakker, F. Dikker, F. Tempelman, and P. Wognum. Diagnosing and solving over-determined constraint satisfaction problems. In *IJCAI*, pages 276–281, 1993.
- BG01. L. Bachmair and H. Ganzinger. Resolution theorem proving. In J. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, pages 19–99. Elsevier and MIT Press, 2001.
- BGJ⁺07. R. Bloem, S. Galler, B. Jobstmann, N. Piterman, A. Pnueli, and M. Weiglhofer. Specify, compile, run: Hardware from PSL. In S. Glesner, J. Knoop, and R. Drechsler, editors, *COCV*, volume 190(4) of *ENTCS*, pages 3–16. Elsevier, 2007.
- BJ. A. Biere and T. Jussila. Benchmark tool run. <http://fmv.jku.at/run/>.
- boo. <http://www.boost.org/doc/libs/release/libs/graph/>.
- CGH97. E. Clarke, O. Grumberg, and K. Hamaguchi. Another look at LTL model checking. *Formal Methods in System Design*, 10(1):47–71, 1997.
- CGP01. E. Clarke, O. Grumberg, and D. Peled. *Model checking*. MIT Press, 2001.
- CMT11. A. Cimatti, S. Mover, and S. Tonetta. Proving and explaining the unfeasibility of message sequence charts for hybrid systems. In P. Bjesse and A. Slobodová, editors, *FMCAD*, pages 54–62, 2011.
- CRST07. A. Cimatti, M. Roveri, V. Schuppan, and S. Tonetta. Boolean abstraction for temporal logic satisfiability. In W. Damm and H. Hermanns, editors, *CAV*, volume 4590 of *Lecture Notes in Computer Science*, pages 532–546. Springer, 2007.
- CRST08. A. Cimatti, M. Roveri, V. Schuppan, and A. Tchaltsev. Diagnostic information for realizability. In F. Logozzo, D. Peled, and L. Zuck, editors, *VMCAI*, volume 4905 of *Lecture Notes in Computer Science*, pages 52–67. Springer, 2008.

- CTVW03. E. Clarke, M. Talupur, H. Veith, and D. Wang. SAT based predicate abstraction for hardware verification. In E. Giunchiglia and A. Tacchella, editors, *SAT*, volume 2919 of *Lecture Notes in Computer Science*, pages 78–92. Springer, 2003.
- Dix95. C. Dixon. *Strategies for Temporal Resolution*. PhD thesis, Department of Computer Science, University of Manchester, 1995. Available from <ftp://ftp.cs.man.ac.uk/pub/TR/UMCS-95-12-1.ps.Z>.
- Dix96. C. Dixon. Search strategies for resolution in temporal logics. In M. McRobbie and J. Slaney, editors, *CADE*, volume 1104 of *Lecture Notes in Computer Science*, pages 673–687. Springer, 1996.
- Dix97. C. Dixon. Using Otter for temporal resolution. In H. Barringer, M. Fisher, D. Gabbay, and G. Gough, editors, *ICTL*, Applied Logic Series, pages 149–166. Kluwer, 1997.
- Dix98. C. Dixon. Temporal resolution using a breadth-first search algorithm. *Ann. Math. Artif. Intell.*, 22(1-2):87–115, 1998.
- EF06. C. Eisner and D. Fisman. *A Practical Introduction to PSL*. Springer, 2006.
- EL86. E. Emerson and C. Lei. Efficient model checking in fragments of the propositional mu-calculus (extended abstract). In *LICS*, pages 267–278. IEEE Computer Society, 1986.
- Eme90. E. Emerson. Temporal and modal logic. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics (B)*, pages 995–1072. 1990.
- FDP01. M. Fisher, C. Dixon, and M. Peim. Clausal temporal resolution. *ACM Trans. Comput. Log.*, 2(1):12–56, 2001.
- FFK⁺01. K. Fisler, R. Fraer, G. Kamhi, M. Vardi, and Z. Yang. Is there a best symbolic cycle-detection algorithm? In T. Margaria and W. Yi, editors, *TACAS*, volume 2031 of *Lecture Notes in Computer Science*, pages 420–434. Springer, 2001.
- Fis91. M. Fisher. A resolution method for temporal logic. In *IJCAI*, pages 99–104, 1991.
- FKSFV08. D. Fisman, O. Kupferman, S. Sheinvald-Faragy, and M. Vardi. A framework for inherent vacuity. In H. Chockler and A. Hu, editors, *Haifa Verification Conference*, volume 5394 of *Lecture Notes in Computer Science*, pages 7–22. Springer, 2008.
- FM09. J. Franco and J. Martin. A history of satisfiability. In A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors, *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, pages 3–74. IOS Press, 2009.
- FN92. M. Fisher and P. Noël. Transformation and synthesis in METATEM. Part I: Propositional METATEM. Technical Report UMCS-92-2-1, University of Manchester, Department of Computer Science, 1992. Available from <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.30.4998>.
- GN03. E. Goldberg and Y. Novikov. Verification of proofs of unsatisfiability for CNF formulas. In *DATE*, pages 10886–10891. IEEE Computer Society, 2003.
- Har05. A. Harding. *Symbolic Strategy Synthesis For Games With LTL Winning Conditions*. PhD thesis, University of Birmingham, 2005.
- HH11. F. Hantry and M. Hacid. Handling conflicts in depth-first search for ltl tableau to debug compliance based languages. In E. Pimentel and V. Valero, editors, *FLACOS*, volume 68 of *EPTCS*, pages 39–53, 2011.

- HK03. U. Hustadt and B. Konev. TRP++ 2.0: A temporal resolution prover. In F. Baader, editor, *CADE*, volume 2741 of *Lecture Notes in Computer Science*, pages 274–278. Springer, 2003.
- HK04. U. Hustadt and B. Konev. TRP++: A temporal resolution prover. In M. Baaz, J. Makowsky, and A. Voronkov, editors, *Collegium Logicum*, volume 8, pages 65–79. Kurt Gödel Society, 2004.
- HKQ03. T. Henzinger, O. Kupferman, and S. Qadeer. From pre-historic to post-modern symbolic model checking. *Formal Methods in System Design*, 23(3):303–327, 2003.
- HKSV01. R. Hardin, R. Kurshan, S. Shukla, and M. Vardi. A new heuristic for bad cycle detection using bdds. *Formal Methods in System Design*, 18(2):131–140, 2001.
- HS02. U. Hustadt and R. A. Schmidt. Scientific benchmarking with temporal logic decision procedures. In D. Fensel, F. Giunchiglia, D. McGuinness, and M. Williams, editors, *KR*, pages 533–546. Morgan Kaufmann, 2002.
- HTKB92. R. Hojati, H. Touati, R. Kurshan, and R. Brayton. Efficient ω -regular language containment. In G v. Bochmann and D. Probst, editors, *CAV*, volume 663 of *Lecture Notes in Computer Science*, pages 396–409. Springer, 1992.
- JB06. B. Jobstmann and R. Bloem. Optimizations for LTL synthesis. In *FM-CAD*, pages 117–124. IEEE Computer Society, 2006.
- JSB06. T. Jussila, C. Sinz, and A. Biere. Extended resolution proofs for symbolic sat solving with quantification. In A. Biere and C. Gomes, editors, *SAT*, volume 4121 of *Lecture Notes in Computer Science*, pages 54–60. Springer, 2006.
- KHB09. R. Könighofer, G. Hofferek, and R. Bloem. Debugging formal specifications using simple counterstrategies. In *FM-CAD*, pages 152–159. IEEE, 2009.
- KHB10. R. Könighofer, G. Hofferek, and R. Bloem. Debugging unrealizable specifications with model-based diagnosis. In S. Barner, I. Harris, D. Kroening, and O. Raz, editors, *HVC*, volume 6504 of *Lecture Notes in Computer Science*, pages 29–45. Springer, 2010.
- Kup06. O. Kupferman. Sanity checks in formal verification. In C. Baier and H. Hermanns, editors, *CONCUR*, volume 4137 of *Lecture Notes in Computer Science*, pages 37–51. Springer, 2006.
- KV03. O. Kupferman and M. Vardi. Vacuity detection in temporal model checking. *STTT*, 4(2):224–233, 2003.
- LH10. M. Ludwig and U. Hustadt. Implementing a fair monodic temporal logic prover. *AI Commun.*, 23(2-3):69–96, 2010.
- Noë95. P. Noë. A transformation-based synthesis of temporal specifications. *Formal Asp. Comput.*, 7(6):587–619, 1995.
- pap. <http://www.schuppan.de/viktor/nfm13/>.
- Pnu77. A. Pnueli. The temporal logic of programs. In *FOCS*, pages 46–57. IEEE, 1977.
- PSC⁺06. I. Pill, S. Semprini, R. Cavada, M. Roveri, R. Bloem, and A. Cimatti. Formal analysis of hardware requirements. In E. Sentovich, editor, *DAC*, pages 821–826. ACM, 2006.
- RBS00. K. Ravi, R. Bloem, and F. Somenzi. A comparative study of symbolic algorithms for the computation of fair cycles. In W. Hunt and S. Johnson, editors, *FM-CAD*, volume 1954 of *Lecture Notes in Computer Science*, pages 143–160. Springer, 2000.

- RKG11. V. Raman and H. Kress-Gazit. Analyzing unsynthesizable specifications for high-level robot behavior using LTLMoP. In G. Gopalakrishnan and S. Qadeer, editors, *CAV*, volume 6806 of *Lecture Notes in Computer Science*, pages 663–668. Springer, 2011.
- Rob65. J. Robinson. A machine-oriented logic based on the resolution principle. *J. ACM*, 12(1):23–41, 1965.
- RV10. K. Rozier and M. Vardi. LTL satisfiability checking. *STTT*, 12(2):123–137, 2010.
- Sch12a. V. Schuppan. Extracting unsatisfiable cores for LTL via temporal resolution (full version), 2012. <http://www.schuppan.de/viktor/nfm13/VSchuppan-NFM-2013-full.pdf>.
- Sch12b. V. Schuppan. Towards a notion of unsatisfiable and unrealizable cores for LTL. *Sci. Comput. Program.*, 77(7-8):908–939, 2012.
- SD11. V. Schuppan and L. Darmawan. Evaluating LTL satisfiability solvers. In T. Bultan and P. Hsiung, editors, *ATVA*, volume 6996 of *Lecture Notes in Computer Science*, pages 397–413. Springer, 2011.
- Sil10. J. Marques Silva. Minimal unsatisfiability: Models, algorithms and applications (invited paper). In *ISMVL*, pages 9–14. IEEE Computer Society, 2010.
- trp. <http://www.csc.liv.ac.uk/~konev/software/trp++/>.
- WDMR08. M. De Wulf, L. Doyen, N. Maquet, and J. Raskin. Antichains: Alternative algorithms for LTL satisfiability and model-checking. In C. Ramakrishnan and J. Rehof, editors, *TACAS*, volume 4963 of *Lecture Notes in Computer Science*, pages 63–77. Springer, 2008.
- ZM03a. L. Zhang and S. Malik. Validating SAT solvers using an independent resolution-based checker: Practical implementations and other applications. In *DATE*, pages 10880–10885. IEEE Computer Society, 2003.
- ZM03b. L. Zhang and S. Malik. Extracting small unsatisfiable cores from unsatisfiable Boolean formula. Presented at *Theory and Applications of Satisfiability Testing, 6th International Conference, SAT 2003. Santa Margherita Ligure, Italy, May 5-8, 2003*.

A TR — A High Level View

In Alg. 2 we turn the discussion in Sec. 3.3 into a high level view of TR in TRP++. At the right hand side of Alg. 2 we show the corresponding line(s) in Alg. 1. When we write “Restrict G to the set of vertices V' ”, we mean that V is intersected with V' , E with $V' \times V'$, and I with V' .

Algorithm 2: High level view of LTL satisfiability checking via TR in TRP++.

Input: A set of SNF clauses C .
Output: *Unsat* if C is unsatisfiable; *sat* otherwise.

```

1  if  $C$  contains the empty clause then return unsat;           // 1
2  Let  $G = (V, E, I)$  be the transition system induced by  $C$ ;   // -
3  Restrict  $G$  to the set of vertices that are the start of an infinite path; // 2
4  if  $I$  is empty then return unsat;                             // 2
5  for  $c \in C$  .  $c = (\mathbf{G}((P) \vee (\mathbf{F}(l))))$  is an eventuality clause in  $C$  do // 3
6  |   Restrict  $G$  to the set of vertices that fulfill  $(P \vee l \vee wl)$ ; // 3
7  |   Restrict  $E$  to the set of edges that fulfill  $((\neg wl) \vee \mathbf{X}(l \vee wl))$ ; // 3
8  Restrict  $G$  to the set of vertices that are the start of an infinite path; // 4
9  if  $I$  is empty then return unsat;                             // 4
10  $G' \leftarrow (\emptyset, \emptyset, \emptyset)$ ;                     // 5
11 while  $G' \neq G$  do                                           // 6
12 |    $G' \leftarrow G$ ;                                         // 7
13 |   for  $c \in C$  .  $c = (\mathbf{G}((P) \vee (\mathbf{F}(l))))$  is an eventuality clause in  $C$  do // 8
14 | |    $V'' \leftarrow \{v \in V \mid \text{a successor of } v \text{ can reach a vertex } v' \text{ in which } l \text{ holds}\}$ ; // 9-16
15 | |   if  $V'' \neq 2^{AP}$  then                                   // 17
16 | | |   Restrict  $G$  to the set of vertices that fulfill  $(P \vee l \vee V'')$ ; // 18
17 | | |   Restrict  $E$  to the set of edges that fulfill  $((\neg wl) \vee \mathbf{X}(l \vee V''))$ ; // 18
18 | | |   Restrict  $G$  to the set of vertices that are the start of an infinite path; // 19
19 | | |   if  $I$  is empty then return unsat;                     // 19
20 return sat;                                                  // 20

```

In model checking a number of works investigated cycle detection algorithms (e.g., [FFK⁺01, RBS00, HKSV01, HTKB92, EL86]). While the high level algorithm above is not identical to any of those, it is somewhat similar to the variant of the One-Way-Catch-Them-Young (OWCTY) algorithm mentioned in footnote 4 of [FFK⁺01]. In particular, it proceeds in backward direction (e.g., [HKQ03]), it uses CTY style pruning (lines 3, 8, 18; [HKSV01]), and the pruning happens in each iteration over the eventuality clauses (resp. fair sets). However, the initial pruning (line 3 or 8) is not present in that algorithm in [FFK⁺01].

B Proofs: 4 UC Extraction

Lemma 1. *Let C be a set of SNF clauses to which Alg. 1 has been applied and shown unsatisfiability, let G be the resolution graph, and let G' the sub-graph according to Def. 2. Let v be a vertex in G' labeled with a clause $c = (\mathbf{G}((\neg wl) \vee (\mathbf{X}(l \vee wl))))$ created by augmentation $\boxed{\text{aug2}}$ from some eventuality clause $(\mathbf{G}((p_1 \vee \dots \vee p_n) \vee (\mathbf{F}(l)))) \in C$ with eventuality literal l . Then there is a*

vertex v' in G' labeled with an eventuality clause $c' = (\mathbf{G}((q_1 \vee \dots \vee q_n) \vee (\mathbf{F}(l)))) \in C$ with eventuality literal l .

Proof. There exists a path π of non-zero length in G' from v to the unique vertex v_\square in the main partition M labeled with the empty clause \square . On the path π there exist two vertices v'', v''' such that v'' is labeled with a clause c'' that contains $\neg wl$ or $\mathbf{X}\neg wl$, while v''' and all of its successors on π are labeled with clauses that contain neither $\neg wl$ nor $\mathbf{X}\neg wl$. Let c''' be the clause labeling v''' .

Case 1. c''' is generated by initial or step resolution $\boxed{\text{init-ii}}$, $\boxed{\text{init-in}}$, $\boxed{\text{step-nn}}$, $\boxed{\text{step-nx}}$, or $\boxed{\text{step-xx}}$ from c'' and some other clause c'''' . c'''' must contain wl or $\mathbf{X}wl$. Moreover, there must be a path π' (possibly of zero length) that starts from a vertex v'''' labeled with a clause c'''' and that ends in the vertex v''' labeled with c''' , such that each vertex on the path π' is labeled with a clause that contains wl or $\mathbf{X}wl$. Finally, wl or $\mathbf{X}wl$ must be present in c'''' either because c'''' is contained in the set of input clauses in SNF, C , or because c'''' is generated by some production rule that introduces wl or $\mathbf{X}wl$ in the conclusion.

Case 1.1. c'''' is contained in the set of input clauses in SNF, C . Impossible: wl is a fresh proposition in $\boxed{\text{aug1}}$ and $\boxed{\text{aug2}}$.

Case 1.2. c'''' is generated by initial or step resolution $\boxed{\text{init-ii}}$, $\boxed{\text{init-in}}$, $\boxed{\text{step-nn}}$, $\boxed{\text{step-nx}}$, or $\boxed{\text{step-xx}}$. Impossible: initial and step resolution do not generate literals that are not contained (modulo time-shifting) in at least one of the premises.

Case 1.3. c'''' is generated by augmentation 1 $\boxed{\text{aug1}}$. By construction of the resolution graph G and the subgraph G' there is an edge in G' from a vertex v' in G' labeled with an eventuality clause $c' = (\mathbf{G}((q_1 \vee \dots \vee q_n) \vee (\mathbf{F}(l)))) \in C$ with eventuality literal l to v'''' .

Case 1.4. c'''' is generated by augmentation 2 $\boxed{\text{aug2}}$, i.e., $c'''' = c$. This introduces another occurrence of $\neg wl$ to be “resolved away”. Note that in the main partition only new clauses are generated from existing ones with edges leading from existing vertices labeled with existing clauses to new vertices labeled with new clauses. Therefore, the main partition of G' is a finite directed acyclic graph, and this case cannot happen infinitely often.

Case 1.5. c'''' is generated by BFS loop search initialization $\boxed{\text{BFS-loop-it-init-x}}$. Impossible: the production rule $\boxed{\text{BFS-loop-it-init-x}}$ copies a clause verbatim. I.e., it cannot be the case that c'''' contains wl or $\mathbf{X}wl$, while the premise does not.

Case 1.6. c'''' is generated by BFS loop search initialization $\boxed{\text{BFS-loop-it-init-n}}$. Impossible: the production rule $\boxed{\text{BFS-loop-it-init-n}}$ copies and time-shifts a clause. I.e., it cannot be the case that c'''' contains $\mathbf{X}wl$, while the premise does not contain wl .

Case 1.7. c'''' is generated by BFS loop search initialization $\boxed{\text{BFS-loop-it-init-c}}$. Impossible: the production rule $\boxed{\text{BFS-loop-it-init-c}}$ copies and time-shifts a clause from a previous BFS loop search

iteration (or initializes with the empty clause \square) and disjoins with an eventuality literal $\mathbf{X}l'$. I.e., it cannot be the case that c'''' contains $\mathbf{X}wl$, while the premise does not contain wl .

Case 1.8. v'''' is linked to via BFS loop search subsumption $\boxed{\text{BFS-loop-it-sub}}$.

This case can be ignored as BFS loop search subsumption $\boxed{\text{BFS-loop-it-sub}}$ does not actually generate a clause but merely links existing ones.

Case 1.9. c'''' is generated by BFS loop search conclusion 1

$\boxed{\text{BFS-loop-conclusion1}}$. Impossible: production rule $\boxed{\text{BFS-loop-conclusion1}}$ copies all literals verbatim from a clause derived in loop search, copies all literals verbatim from an eventuality clause except for the eventuality literal l' prefixed by \mathbf{F} , and disjoins with the eventuality literal l' . I.e., it cannot be the case that c'''' contains wl , while the premises do not.

Case 1.10. c'''' is generated by BFS loop search conclusion 2

$\boxed{\text{BFS-loop-conclusion2}}$. Impossible: production rule $\boxed{\text{BFS-loop-conclusion2}}$ copies and time-shifts all literals from a clause c'''' derived in loop search and disjoins with $\neg wl'$ and $\mathbf{X}l'$ for some eventuality literal l' . I.e., it cannot be the case that c'''' contains $\mathbf{X}wl$, while the premise c'''' does not contain wl .

Case 2. c''' is generated by augmentation $\boxed{\text{aug1}}$ or $\boxed{\text{aug2}}$. Impossible: the premise of the production rules $\boxed{\text{aug1}}$ and $\boxed{\text{aug2}}$ cannot contain either $\neg wl$ or $\mathbf{X}\neg wl$ as wl is assumed to be a fresh proposition in $\boxed{\text{aug1}}$ and $\boxed{\text{aug2}}$.

Case 3. c''' is generated by BFS loop search initialization $\boxed{\text{BFS-loop-it-init-x}}$. Impossible: the production rule $\boxed{\text{BFS-loop-it-init-x}}$ copies a clause verbatim. I.e., it cannot be the case that c''' contains $\neg wl$ or $\mathbf{X}\neg wl$, while c''' does not.

Case 4. c''' is generated by BFS loop search initialization $\boxed{\text{BFS-loop-it-init-n}}$. Impossible: the production rule $\boxed{\text{BFS-loop-it-init-n}}$ copies and time-shifts a clause. I.e., it cannot be the case that c''' contains $\neg wl$, while c''' does not contain $\mathbf{X}\neg wl$.

Case 5. c''' is generated by BFS loop search initialization $\boxed{\text{BFS-loop-it-init-c}}$. Impossible: the production rule $\boxed{\text{BFS-loop-it-init-c}}$ copies and time-shifts a clause from a previous BFS loop search iteration (or initializes with the empty clause \square) and disjoins with an eventuality literal $\mathbf{X}l'$. I.e., it cannot be the case that c''' contains $\neg wl$, while c''' does not contain $\mathbf{X}\neg wl$.

Case 6. v'' and v''' are linked via BFS loop search subsumption $\boxed{\text{BFS-loop-it-sub}}$, i.e., a time-shifted version of c'' subsumes c''' . Impossible: $\boxed{\text{BFS-loop-it-sub}}$ links from a clause with fewer literals to a clause with (modulo time-shifting) the same and more literals. I.e., it cannot be the case that c'' contains $\neg wl$, while c''' does not contain $\mathbf{X}\neg wl$.

Case 7. c''' is generated by BFS loop search conclusion 1 $\boxed{\text{BFS-loop-conclusion1}}$.

Impossible: production rule $\boxed{\text{BFS-loop-conclusion1}}$ copies all literals verbatim from a clause derived in loop search, copies all literals verbatim from an eventuality clause except for the eventuality literal l' prefixed by \mathbf{F} , and disjoins with the eventuality literal l' . I.e., it cannot be the case that c''' contains $\neg wl$, while c''' does not.

Case 8. c''' is generated by BFS loop search conclusion 2 BFS-loop-conclusion2.

Impossible: production rule BFS-loop-conclusion2 copies and time-shifts all literals from a clause derived in loop search and disjoins with $\neg wl'$ and $\mathbf{X}l'$ for some eventuality literal l' . I.e., it cannot be the case that c'' contains $\neg wl$, while c''' does not contain $\mathbf{X}\neg wl$.

Notice that the only possible cases are case 1.3 and 1.4. Of those, case 1.4 can only happen a finite number of times and must be followed by an occurrence of case 1.3. This concludes the proof.

Lemma 2. *Let C be a set of SNF clauses to which Alg. 1 has been applied and shown unsatisfiability, let G be the resolution graph constructed, and let G' be the subgraph according to Def. 2. Let v be a vertex in G' labeled with a clause $c = (\mathbf{G}((\neg wl) \vee (\mathbf{X}((q_1 \vee \dots \vee q_{n'}) \vee l))))$ generated by BFS loop search conclusion 2 BFS-loop-conclusion2 from some eventuality clause $(\mathbf{G}((p_1 \vee \dots \vee p_n) \vee (\mathbf{F}(l)))) \in C$ with eventuality literal l (and some other clause). Then there is a vertex v'' in G' labeled with an eventuality clause $c'' = (\mathbf{G}((r_1 \vee \dots \vee r_{n''}) \vee (\mathbf{F}(l)))) \in C$ with eventuality literal l .*

Proof. Analogous to the proof of Lemma 1.

Lemma 3. *Let C be a set of SNF clauses to which Alg. 1 has been applied and shown unsatisfiability, let G be the resolution graph, and let G' be the subgraph according to Def. 2. Let v be a vertex in G' labeled with a clause $c = (\mathbf{G}((0) \vee (\mathbf{X}(q_1 \vee \dots \vee q_{n'} \vee l))))$ generated by production rule BFS-loop-it-init-c from some eventuality clause $(\mathbf{G}((p_1 \vee \dots \vee p_n) \vee (\mathbf{F}(l)))) \in C$ with eventuality literal l (and some other clause). Then there is a vertex v'' in G' labeled with an eventuality clause $c'' = (\mathbf{G}((r_1 \vee \dots \vee r_{n''}) \vee (\mathbf{F}(l)))) \in C$ with eventuality literal l .*

Proof. By construction of the resolution graph G (Def. 1) and its subgraph G' (Def. 2) v is included in G' only if G' also includes some vertex v' labeled with some clause c' such that c' was generated by BFS loop search conclusion BFS-loop-conclusion1 or BFS-loop-conclusion2 from the BFS loop search iteration of which c is part.

Case 1. c' is generated by BFS loop search conclusion 1 BFS-loop-conclusion1.

The claim follows from the construction of the resolution graph G and its subgraph G' . By Def. 1 v' has an incoming edge from a vertex v'' labeled with an eventuality clause $c'' = (\mathbf{G}((r_1 \vee \dots \vee r_{n''}) \vee (\mathbf{F}(l)))) \in C$ with eventuality literal l and by Def. 2 v'' is included in G' if v' is included.

Case 2. c' is generated by BFS loop search conclusion 2 BFS-loop-conclusion2. In that case the claim follows directly from Lemma 2.

Theorem 1 (Unsatisfiability of UC in SNF). *Let C be a set of SNF clauses to which Alg. 1 has been applied and shown unsatisfiability, and let C^{uc} be the UC of C in SNF. Then C^{uc} is unsat.*

Proof. Assume for a moment that in columns 8 ($p.1 - c$) and 9 ($p.2 - c$) of Tab. 1 all \star are replaced with \checkmark , i.e., that each conclusion in the resolution graph is connected by an edge to each of its premises rather than only to a subset of them. In that case the UC in SNF according to Def. 2 would contain all clauses of the set of starting clauses C that contributed to deriving the empty clause and, hence, to establishing unsatisfiability of C . In that case it would follow directly from the correctness of TR that C^{uc} is unsatisfiable.

It remains to show that 1. not including an edge from premise 1 to the conclusion for rule $\boxed{\text{aug2}}$, 2. not including an edge from premise 2 to the conclusion for rule $\boxed{\text{BFS-loop-conclusion2}}$, 3. not including an edge from premise 2 to the conclusion for rule $\boxed{\text{BFS-loop-it-init-c}}$, and 4. not including an edge from premise 1 to the conclusion for rule $\boxed{\text{BFS-loop-it-init-c}}$ in the resolution graph G maintains the fact that the resulting C^{uc} is unsatisfiable. Items 1. – 3. are addressed by Lemmas 1, 2, and 3.

We now address item 4. Notice that this case essentially corresponds to considering only the last iteration of a successful loop search to obtain the UC C^{uc} . After initialization of a loop search iteration in line 11 of Alg. 1 L contains three sets of clauses according to the three production rules for initializing a loop search iteration. Clauses generated by $\boxed{\text{BFS-loop-it-init-x}}$ and $\boxed{\text{BFS-loop-it-init-n}}$ are (partly time-shifted) duplicates of clauses derived so far in the main partition. $\boxed{\text{BFS-loop-it-init-c}}$ generates a set of clauses $(\mathbf{G}((0) \vee (\mathbf{X}(p_{i,1} \vee \dots \vee p_{i,n_i} \vee l))))$. From these three sets saturation restricted to rule $\boxed{\text{step-xx}}$ in line 12 derives another set of clauses $(\mathbf{G}(q_{i',1} \vee \dots \vee q_{i',n'_{i'}}))$. Taking the restriction of saturation to rule $\boxed{\text{step-xx}}$ into account, that loop search iteration has established that, assuming C , the following fact is provable:

$$\mathbf{G}((\bigwedge_{1 \leq i \leq n} (\mathbf{X}(p_{i,1} \vee \dots \vee p_{i,n_i} \vee l))) \rightarrow (\bigwedge_{1 \leq i' \leq n'} (q_{i',1} \vee \dots \vee q_{i',n'_{i'}}))) \quad (10)$$

Moreover, if subsumption in line 15 succeeds, the following fact is also provable:

$$\bigwedge_{1 \leq i \leq n} (\bigvee_{1 \leq i' \leq n'} (\mathbf{G}((q_{i',1} \vee \dots \vee q_{i',n'_{i'}}) \rightarrow (p_{i,1} \vee \dots \vee p_{i,n_i})))) \quad (11)$$

We rewrite (10) and (11) as follows:

$$\begin{aligned} & \mathbf{G}((\bigwedge_{1 \leq i \leq n} (\mathbf{X}(p_{i,1} \vee \dots \vee p_{i,n_i} \vee l))) \rightarrow (\bigwedge_{1 \leq i' \leq n'} (q_{i',1} \vee \dots \vee q_{i',n'_{i'}}))) \\ \Leftrightarrow & \mathbf{G}((\bigwedge_{1 \leq i' \leq n'} ((\bigwedge_{1 \leq i \leq n} (\mathbf{X}(p_{i,1} \vee \dots \vee p_{i,n_i} \vee l))) \rightarrow (q_{i',1} \vee \dots \vee q_{i',n'_{i'}})))) \\ \Leftrightarrow & \bigwedge_{1 \leq i' \leq n'} (\mathbf{G}((\bigwedge_{1 \leq i \leq n} (\mathbf{X}(p_{i,1} \vee \dots \vee p_{i,n_i} \vee l))) \rightarrow (q_{i',1} \vee \dots \vee q_{i',n'_{i'}}))) \\ \Leftrightarrow & \bigwedge_{1 \leq i' \leq n'} (\mathbf{G}((\neg(q_{i',1} \vee \dots \vee q_{i',n'_{i'}})) \rightarrow (\neg(\bigwedge_{1 \leq i \leq n} (\mathbf{X}(p_{i,1} \vee \dots \vee p_{i,n_i} \vee l)))))) \\ \Leftrightarrow & \bigwedge_{1 \leq i' \leq n'} (\mathbf{G}((\neg(q_{i',1} \vee \dots \vee q_{i',n'_{i'}})) \rightarrow (\bigvee_{1 \leq i \leq n} (\mathbf{X}(\neg(p_{i,1} \vee \dots \vee p_{i,n_i} \vee l)))))) \\ \Leftrightarrow & \bigwedge_{1 \leq i' \leq n'} (\mathbf{G}((\neg(q_{i',1} \vee \dots \vee q_{i',n'_{i'}})) \rightarrow (\bigvee_{1 \leq i \leq n} (\mathbf{X}((\neg(p_{i,1} \vee \dots \vee p_{i,n_i})) \wedge (\neg l)))))) \\ \Leftrightarrow & \bigwedge_{1 \leq i' \leq n'} (\mathbf{G}((\neg(q_{i',1} \vee \dots \vee q_{i',n'_{i'}})) \rightarrow ((\mathbf{X} \neg l) \wedge (\bigvee_{1 \leq i \leq n} (\mathbf{X}(\neg(p_{i,1} \vee \dots \vee p_{i,n_i}))))))) \quad (12) \end{aligned}$$

$$\begin{aligned}
& \bigwedge_{1 \leq i \leq n} \left(\bigvee_{1 \leq i' \leq n'} (\mathbf{G}((q_{i',1} \vee \dots \vee q_{i',n'}_{i'}) \rightarrow (p_{i,1} \vee \dots \vee p_{i,n_i}))) \right) \\
& \Leftrightarrow \bigwedge_{1 \leq i \leq n} \left(\bigvee_{1 \leq i' \leq n'} (\mathbf{G}((\neg(p_{i,1} \vee \dots \vee p_{i,n_i})) \rightarrow (\neg(q_{i',1} \vee \dots \vee q_{i',n'}_{i'})))) \right) \quad (13)
\end{aligned}$$

Putting (12) and (13) together, we obtain (14), which is exactly the premise required to perform eventuality resolution with an eventuality clause with eventuality literal l [FDP01]:

$$\begin{aligned}
& (\mathbf{G}((q_{1,1} \vee \dots \vee q_{1,n'}_1) \vee (\mathbf{XG}\neg l))) \\
& \quad \dots \\
& (\mathbf{G}((q_{n',1} \vee \dots \vee q_{n',n'}_{n'}) \vee (\mathbf{XG}\neg l))) \quad (14)
\end{aligned}$$

This concludes the proof.

Proposition 1 (Complexity of UC Extraction). *Let C be a set of SNF clauses to which Alg. 1 is applied and shows unsatisfiability. Construction and backward traversal of the resolution graph and, hence, construction of C^{uc} according to Def. 2 can be performed in time $\mathcal{O}(|V|)$ in addition to the time required to run Alg. 1. $|V|$ is at most exponential in $|AP| + \log(|C|)$.*

Proof. Notice that each vertex in G has at most 2 incoming edges. Hence, construction of G and backward traversal of G from the unique vertex in the main partition labeled with the empty clause, v_\square , can be performed in time $\mathcal{O}(|V|)$.

For a proof of $|AP| + \log(|C|)$ see the following reasoning:

1. In an initial clause a proposition can be not present, present, or present negated. Hence, the number of different initial clauses is $\mathcal{O}(3^{|AP|})$.
2. In a global clause a proposition can be one of not present, present, or present negated; and prefixed by \mathbf{X} not present, present, or present negated. Hence, the number of different global clauses is $\mathcal{O}(9^{|AP|})$.
3. The number of clauses in the main partition is bounded by $|C| + \mathcal{O}(3^{|AP|}) + \mathcal{O}(9^{|AP|}) = \mathcal{O}(|C| + 9^{|AP|})$.
4. The number of clauses in a partition for a BFS loop search iteration is bounded by $\mathcal{O}(9^{|AP|})$.
5. The number of partitions is bounded by 1 plus the number of BFS loop search iterations.
6. The number of iterations in a BFS loop search is bounded by the length of the longest monotonically increasing sequence of Boolean formulas over AP , which is $\mathcal{O}(2^{|AP|})$. See also [Dix98].
7. The number of BFS loop searches is bounded by the number of different clauses that can be the result of a BFS loop search. The number of different clauses that can be the consequence of BFS loop search conclusion 1 BFS-loop-conclusion1 is bounded by the number of different global clauses with empty next part, which is $\mathcal{O}(3^{|AP|})$. The number of different clauses that can be the consequence of BFS loop search conclusion 2 BFS-loop-conclusion2 is

- bounded by the number of different eventuality literals times the number of different global clauses with empty next part, which is $\mathcal{O}(|C| \cdot 3^{|AP|})$. Hence, the number of BFS loop searches is bounded by $\mathcal{O}(|C| \cdot 3^{|AP|})$.
8. Taking all of the above into account, the number of clauses is bounded by $\mathcal{O}(|C| + 9^{|AP|} + |C| \cdot 3^{|AP|} \cdot 2^{|AP|} \cdot 9^{|AP|}) = \mathcal{O}(|C| \cdot 54^{|AP|})$.

This concludes the proof.

C Proofs: 5 From LTL to SNF and Back

Theorem 2 (Unsatisfiability of UC in LTL). *Let ϕ be an unsatisfiable LTL formula, and let ϕ^{uc} be the UC of ϕ in LTL. Then ϕ^{uc} is unsat.*

Proof. Let $SNF(\phi)$ be the SNF of ϕ , and let C^{uc} be the UC of $SNF(\phi)$ in SNF.

First, consider the trivial case that ϕ is 0. Here, Def. 4 results in the UC of ϕ in LTL being $\phi^{uc} \equiv 0$ as desired.

Now assume that ϕ is not 0, i.e., the size of the syntax tree of ϕ is greater than 1. Let $SNF(\phi^{uc})$ be the SNF of ϕ^{uc} . In order to prove that ϕ^{uc} is *unsat* we show that the clauses of C^{uc} (which is *unsat*) are a subset of the SNF of ϕ^{uc} : $C^{uc} \subseteq SNF(\phi^{uc})$.

By comparing the clauses of $SNF(\phi)$ with those of $SNF(\phi^{uc})$ we can partition the clauses of $SNF(\phi)$ into 3 sets:² 1. Some clauses are present in both $SNF(\phi)$ and $SNF(\phi^{uc})$: $C'_1 \equiv SNF(\phi) \cap SNF(\phi^{uc})$. 2. Some clauses are present in $SNF(\phi)$ and are present in $SNF(\phi^{uc})$ with one or more occurrences of some propositions x, x', \dots that are marked blue boxed in Tab. 2 replaced with 1 or 0. Call that set C'_2 . 3. Some clauses are present in $SNF(\phi)$ but not in $SNF(\phi^{uc})$: $C'_3 \equiv SNF(\phi) \setminus (SNF(\phi^{uc}) \cup C'_2)$.

By Def. 2 C^{uc} is a subset of $SNF(\phi)$: $C^{uc} \subseteq SNF(\phi)$. By Def. 4 C^{uc} contains no member of C'_2 ; otherwise, there could not be one or more occurrences of some propositions x, x', \dots that are marked blue boxed in Tab. 2 replaced with 1 or 0 in the clauses of C'_2 : $C^{uc} \cap C'_2 = \emptyset$. Now we argue that C^{uc} also contains no member of C'_3 . First, let $c \in C'_3$ be an initial or a global clause. c cannot be a member of C^{uc} as, in order to be part of a proof that derives the empty clause, all literals of c need to be “resolved away”. However, this is not possible for c as for the literal $(\neg)x_\psi$ on the left side of the implication in Tab. 2 there is no clause with an opposite literal in C^{uc} . This follows by induction on the nesting depth of the subformula ψ to which $(\neg)x_\psi$ belongs from the occurrence of the superformula of ψ that has been replaced with 1 or 0 in ϕ^{uc} . Now let $c \in C'_3$ be an eventuality clause. By Def. 1, 2 for such c to be part of C^{uc} there would have to be a clause c' in the resolution graph G according to Def. 1 that was generated by production rules aug1 or BFS-loop-conclusion1 and that is backward reachable in G from the vertex labeled with the empty clause \square in the main partition M , v_\square . Again, for the latter to happen, all literals of c' would have to be “resolved away”, which is impossible by a similar inductive argument as

² We disregard the issue of the indices of the variables x, x', \dots

before. Hence, we have shown that all clauses in C^{uc} come from C'_1 , which is a subset of $SNF(\phi^{uc})$. This concludes the proof.

D Additional Plots

Figures 3 and 4 show the overhead that is incurred and the size reduction that is obtained by extracting UCs split by category.

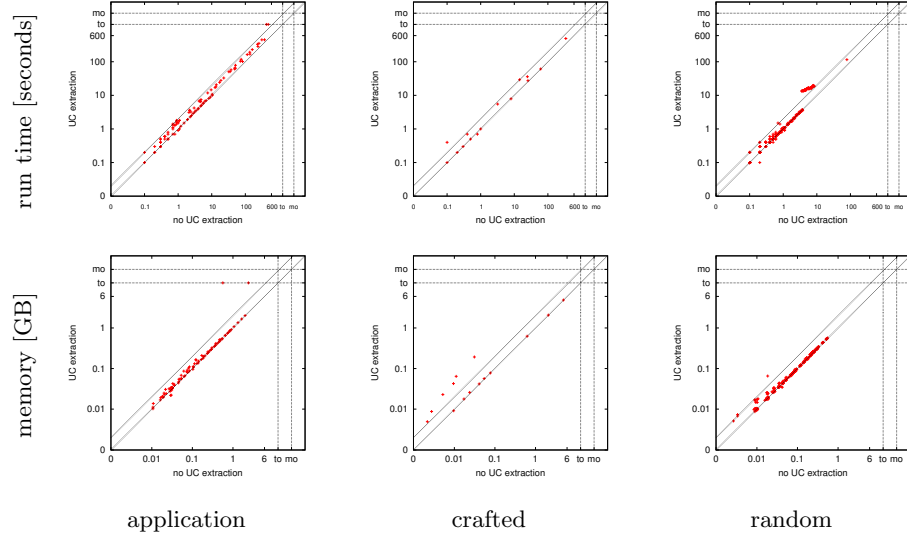


Fig. 3. Overhead incurred by UC extraction in terms of run time (in seconds) and memory (in GB) separated by categories **application**, **crafted**, and **random**. In each graph extraction of UCs is on the y-axis and no UC extraction on the x-axis. The off-center diagonal shows where $y = 2x$.




Fig. 4. Size reduction obtained by UC extraction separated by categories **application**, **crafted**, and **random**. The y-axes show the sizes of the UCs, the x-axes show the sizes of the input formulas. Size is measured as the number of nodes in the syntax trees.